arXiv:1404.7789v1 [cs.SI] 30 Apr 2014

Phase transitions in semisupervised clustering of sparse networks

Pan Zhang[1], Cristopher Moore[1], and Lenka Zdeborová[2]
[1)] *Santa Fe Institute, Santa Fe, New Mexico 87501, USA*
[2)] *Institut de Physique Théorique, CEA Saclay and URA 2306, CNRS, Gif-sur-Yvette, France.*

Predicting labels of nodes in a network, such as community memberships or demographic variables, is an important problem with applications in social and biological networks. A recently-discovered phase transition puts fundamental limits on the accuracy of these predictions if we have access only to the network topology. However, if we know the correct labels of some fraction α of the nodes, we can do better. We study the phase diagram of this "semisupervised" learning problem for networks generated by the stochastic block model. We use the cavity method and the associated belief propagation algorithm to study what accuracy can be achieved as a function of α. For $k = 2$ groups, we find that the detectability transition disappears for any $\alpha > 0$, in agreement with previous work. For larger k where a hard but detectable regime exists, we find that the easy/hard transition (the point at which efficient algorithms can do better than chance) becomes a line of transitions where the accuracy jumps discontinuously at a critical value of α. This line ends in a critical point with a second-order transition, beyond which the accuracy is a continuous function of α. We demonstrate qualitatively similar transitions in two real-world networks.

I. INTRODUCTION

Community or module detection, also known as node clustering, is an important task in the study of biological, social and technological networks. Many methods have been proposed to solve this problem, including spectral clustering [1–3]; modularity optimization [4–7]; statistical inference using generative models, such as the stochastic block model [8–10] and a wide variety of other methods, e.g. [6, 11, 12]. See [13] for a review.

It was shown in [8, 9] that for sparse networks generated by the stochastic block model [14], there is a phase transition in community detection. This transition was initially established using the cavity method, or equivalently by analyzing the behavior of belief propagation. It was recently established rigorously in the case of two groups of equal size [15–17]. In this case, below this transition, no algorithm can label the nodes better than chance, or even distinguish the network from an Erdős-Rényi random graph with high probability. In terms of belief propagation, there is a factorized fixed point where every node is equally likely to be in every group, and it becomes globally stable at the transition.

For more than two groups, there is an additional regime where the factorized fixed point is locally stable, but another, more accurate, fixed point is locally stable as well. This regime lies between two spinodal transitions: the easy/hard transition where the factorized fixed point becomes locally unstable, so that efficient algorithms can achieve a high accuracy (also known as the Kesten-Stigum transition or the robust reconstruction threshold) and the transition where the accurate fixed point first appears (also known as the reconstruction threshold). In between these two, there is a first order phase transition, where the Bethe free energy of these two fixed points cross. This is the detectability transition, in the sense that an algorithm that can search exhaustively for fixed points—which would take exponential time—would choose the accurate fixed point above this transition. However, below this transition there are exponentially many competing fixed points, each corresponding to a cluster of assignments, and even an exponential-time algorithm has no way to tell which is the correct one. (Note that, of these three transitions, the detectability transition is the only true thermodynamic phase transition; the others are dynamical.)

In between the first order phase and easy/hard transitions, there is a "hard but detectable" regime where the communities can be identified in principle; if we could perform an exhaustive search, we would choose the accurate fixed point since it has lower free energy. In Bayesian terms, the correct block model has larger total likelihood than an Erdős-Rényi graph. However, the accurate fixed point has a very small basin of attraction, making it exponentially hard to find—unless we have some additional information.

Here we model this additional information as a so-called *semisupervised* learning problem (e.g. [18]) where we are given the true labels of some small fraction α of the nodes. This information shifts the location of these transitions; in essence, it destabilizes the factorized fixed point, and pushes us towards the basin of attraction of the accurate one. As a result, for some values of the block model parameters, there is a discontinuous jump in the accuracy as a function of α. Roughly speaking, for very small α our information is local, consisting of the known nodes and good guesses about nodes in their vicinity: but at a certain α belief propagation causes this information to percolate, giving us high accuracy throughout the network. As we vary the block model parameters, this line terminates at the point where the two spinodals and first order phase transitions all meet. At that critical point there is a second-order phase transition, and beyond that point the accuracy is a continuous function of α.

Semisupervised learning is an important task in machine learning, in settings where hidden variables or labels are expensive and time-consuming to obtain. Semisupervised community detection was studied in several previous papers [19–21]. The conclusion of [19] was that the detectability transition disappears for any $\alpha > 0$. Later [20] suggested that in some cases it survives for more than two groups. However, both these works were based on an approximate (zero temperature replica symmetric) calculation that corresponds to a far-from-optimal algorithm; moreover, it is known to lead to unphysical results in many other models such as graph coloring (which is a special case of the stochastic block model) or random k-SAT [22, 23].

In the present paper we investigate semisupervised community detection using the cavity method and belief propagation, which in sparse graphs is believed to be Bayes-optimal in the limit $n \to \infty$. From a physics point of view, our results settle the question of what exactly happens in the semisupervised setting, including how the reconstruction, detectability, and easy/hard transitions vary as a function of α. From the point of view of mathematics, our calculations provide non-trivial conjectures that we hope will be amenable to rigorous proof.

Our calculations follow the same methodology as those carried out in two other problems:

- *Study of the ideal glass transition by random pinning.* An important property of Bayes-optimal inference is that the true configuration cannot be distinguished from other configurations that are sampled at random from the posterior probability measure. This is why considering a disordered system similar to the one in this paper and fixing the value or position of a small fraction of nodes in a randomly chosen equilibrium configurations is formally the same problem to semisupervised learning. The analysis of systems with pinned particles was done in order to better understand the formation of glasses in [24, 25].

- *Analysis of belief propagation guided decimation.* belief propagation with decimation is a very interesting solver for a wide range of random constraint satisfaction problems. Its performance was analyzed in [26, 27]. If we decimate a fraction α of the variables (i.e., fix them to particular values) this affects the further performance of the algorithm in a way similar to semisupervised learning.

For random k-SAT, a large part of this picture has been made rigorous [28]. Our hope is that similar techniques will apply to our results here. As a first step, very recent work [29] shows that semisupervised learning does indeed allow for partial reconstruction below the detectability threshold for $k > 4$ groups.

The paper is organized as follows. Section II includes definitions and the description of the stochastic block model. In Sections III we consider semisupervised learning in the networks generated by stochastic block model. In Section IV we consider semisupervised learning in two real-world networks, finding transitions in the accuracy at a critical value of α qualitatively similar to our analysis for the block model. We conclude in Section V.

II. THE STOCHASTIC BLOCK MODEL, BELIEF PROPAGATION, AND SEMISUPERVISED LEARNING

The stochastic block model is defined as follows. Nodes are split into k groups, where each group $1 \le a \le k$ contains an expected fraction q_a of the nodes. Edge probabilities are given by a $k \times k$ matrix p. We generate a random network G with n nodes as follows. First, we choose a group assignment $t \in \{1, \dots, k\}^n$ by assigning each node i a label $t_i \in \{1, ..., k\}$ chosen independently with probability q_{t_i}. Between each pair of nodes i and j, we then add an edge between them with probability p_{t_i,t_j}. For now, we assume that the parameters k, q (a vector denoting $\{q_a\}$), and p (a matrix denoting $\{q_{ab}\}$) are known.

The likelihood of generating G given the parameters and the labels is

$$P(G, t|q, p) = \prod_i q_{t_i} \prod_{\langle ij\rangle \in E} p_{t_i,t_j} \prod_{\langle ij\rangle \notin E} (1 - p_{t_i,t_j}) \,, \tag{1}$$

the Gibbs distribution of the labels t, i.e., their posterior distribution given G, can be computed via Bayes' rule,

$$P(t|G, q, p) = \frac{P(G, t|q, p)}{\sum_s P(G, s|q, p)} \,. \tag{2}$$

In this paper, we consider sparse networks where $p_{ab} = c_{ab}/n$ for some constant matrix c. In this case, the marginal probability ψ_a^i that a given node i has label $t_i = a$ can be computed using belief propagation. The idea of BP is to replace these marginals with "messages" $\psi_a^{i\to\ell}$ from i to each of its neighbors ℓ, which are estimates of these marginals based on i's interactions with its other neighbors [30, 31]. We assume that the neighbors of each node are conditionally

independent of each other; equivalently, we ignore the effect of loops. For the stochastic block model, we obtain the following update equations for these messages [8, 9]:

$$\psi_a^{i\to\ell} = \frac{1}{Z_{i\to\ell}}\, q_a \mathrm{e}^{h_a} \prod_{j\in\partial i\setminus\ell} \sum_b c_{ab}\psi_b^{j\to i}\,, \tag{3}$$

Here $Z_{i\to l}$ is a normalization factor and h_a is an adaptive external field that enforces the expected group sizes,

$$h_a = \frac{1}{n}\sum_i \sum_b c_{ab}\psi_b^i, \tag{4}$$

where the marginal probability that $t_i = a$ is given by

$$\psi_a^i = \frac{1}{Z_i}\, q_a \mathrm{e}^{h_a} \prod_{j\in\partial i} \sum_b c_{ab}\psi_b^{j\to i}\,, \tag{5}$$

In the usual setting, we start with random messages, and apply the BP equations (3) until we reach a fixed point. In order to predict the node labels, we assign each node to its most-likely label according to its marginal:

$$\widehat{t}_i = \operatorname*{argmax}_a \psi_a^i\,.$$

Fixed points of the BP equations are stationary points of the Bethe free energy [30], which up to a constant is

$$F_{\mathrm{Bethe}} = \sum_{\langle ij\rangle\in E} \log\left[\sum_{a,b=1}^{k} c_{ab}\psi_a^{i\to j}\psi_b^{j\to i}\right] - \sum_i \log\left[\sum_{a=1}^{k} q_a \mathrm{e}^{h_a}\prod_{j\in\partial i}\sum_{\{b\}} c_{ab}\psi_b^{j\to i}\right]. \tag{6}$$

If there is more than one fixed point, the one with the lowest F_{Bethe} gives an optimal estimate of the marginals, since F_{Bethe} is minus the logarithm of the total likelihood of the block model. However, as we comment above, if the optimal fixed point has a small basin of attraction, then finding it through exhaustive search will take exponential time. Analyzing stability of instability of these fixed points, including the trivial or "factorized" fixed point where $\psi_a^{i\to\ell} = q_a$, leads to the phase transitions in the stochastic block model described in [8, 9].

It is straightforward to adapt the above formalism to the case of semisupervised learning. One uses exactly the same equations except that nodes whose labels have been revealed have fixed messages. If we know that $t_i = a^*$, then for all ℓ we have

$$\psi_a^{i\to\ell} = \delta_{a,a^*}\,.$$

Equivalently, we can define a local external field, replacing the global parameter q with q^i in (3). Then

$$q_a^i = \delta_{a,a^*}\,.$$

In this paper we focus on a widely-studied special case of the stochastic block model, also well-known as the planted partition model, where the groups are of equal size, i.e., ,$q_a = 1/k$, and where c_{ab} takes only two values:

$$c_{ab} = \begin{cases} c_{\mathrm{in}} & \text{if } a = b \\ c_{\mathrm{out}} & \text{if } a \neq b\,. \end{cases} \tag{7}$$

In that case, the average degree is $c = (c_{\mathrm{in}} + (q-1)c_{\mathrm{out}})/q$. It is common to parametrize this model with the ratio $\epsilon = c_{\mathrm{out}}/c_{\mathrm{in}}$. When $\epsilon = 0$ is small, nodes are connected only to others in the same group; at $\epsilon = 1$ the network is an Erdős-Rényi graph, where every pair of nodes is equally likely to be connected.

We assume here that the parameters $k, c_{\mathrm{in}}, c_{\mathrm{out}}$ are known. If they are unknown, inferring them from the graph and partial information about the nodes is an interesting learning problem in its own right. One can estimate them from the set of known edges, i.e., those where both endpoints have known labels: in the sparse case there are $O(\alpha^2 n)$ such edges, or $O(n)$ if the fraction of known labels α is constant. However, for $\alpha \sim 10^{-2}$, say, there are very few such edges until $n \gtrsim 10^5$ or 10^6. Alternately, we can learn the parameters using the expectation-maximization (EM) algorithm of [8, 9], which minimizes the Bethe free energy, or a hybrid method where we initialize EM with parameters estimated from the known edges, if any.

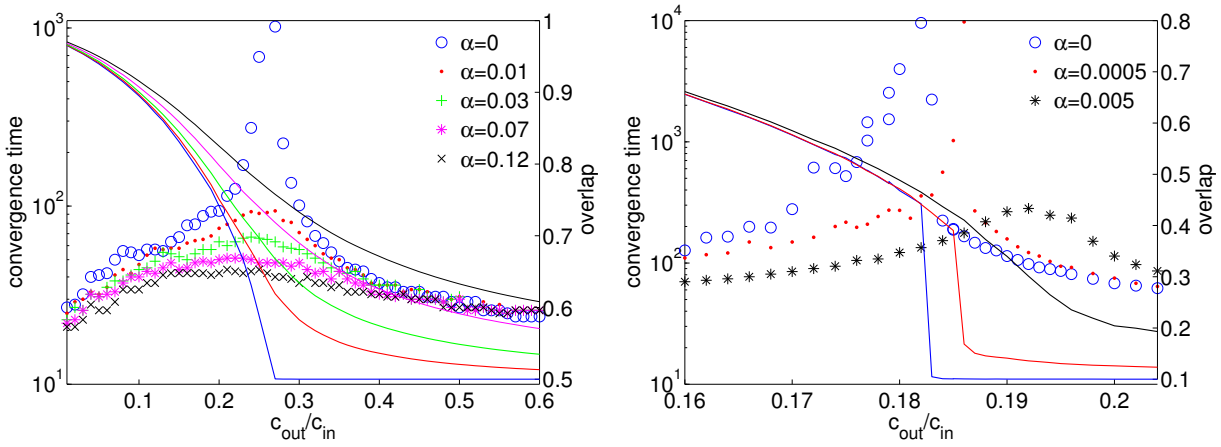

FIG. 1: Overlap and convergence time of BP as a function of $\epsilon = c_{\rm out}/c_{\rm in}$ for different α, on networks generated by the stochastic block model. On the left, $k = 2$, $c = 3$, and $n = 10^5$. For just two groups, the transition disappears for any $\alpha > 0$. On the right, $k = 10$, $c = 10$, $n = 5 \times 10^5$. Here the easy/hard transition persists for small values of α, with a discontinuity in the overlap and a diverging convergence time; this transition disappears at a critical value of α.

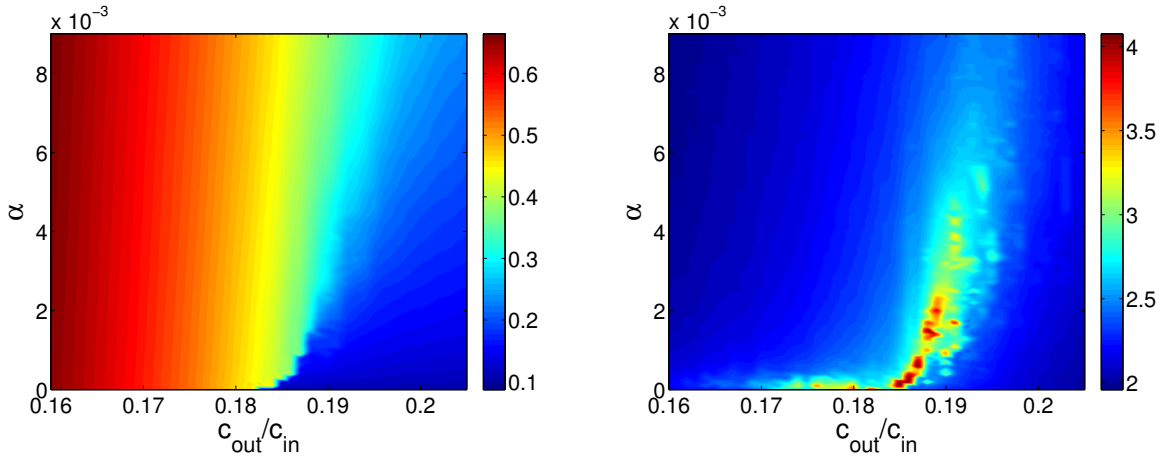

FIG. 2: Left, overlap as a function of $\epsilon = c_{\rm out}/c_{\rm in}$ and α for networks with same parameters as in the right of Fig. 1. The heat map shows a line of discontinuities, ending at a second-order phase transition beyond which the overlap is a smooth function. Right, the logarithm (base 10) of the convergence time in the same plane, showing divergence along the critical line.

III. RESULTS ON THE STOCHASTIC BLOCK MODEL AND THE FATE OF THE TRANSITIONS

First we investigate semisupervised learning for assortative networks, i.e., the case $c_{\rm in} > c_{\rm out}$. As shown in [8], in the unsupervised case $\alpha = 0$ there is a phase transition at

$$c_{\rm in} - c_{\rm out} = k\sqrt{c}\,,$$

where the factorized fixed point goes from stable to unstable. Below this transition the overlap, i.e., the fraction of correctly assigned nodes, is $1/k$, no better than random chance. For $k \leq 4$ this phase transition is second-order: the overlap is continuous, but with discontinuous derivative at the transition. For $k > 4$, it becomes an "easy/hard" transition, with a discontinuity in the overlap when we jump from the factorized fixed point to the accurate one. In both cases, the convergence time (the number of iterations BP takes to reach a fixed point) diverges at the transition.

In Fig. 1 we show the overlap achieved by BP for two different values of k and various values of α. In each case, we hold the average degree c fixed and vary $\epsilon = c_{\rm out}/c_{\rm in}$. On the left, we have $k = 2$. Here, analogous to the unsupervised case $\alpha = 0$, the overlap is a continuous function of ϵ. Moreover, for $\alpha > 0$ the detectability transition disappears: the overlap becomes a smooth function, and the convergence time no longer diverges. This picture agrees qualitatively with the approximate analytical results in [19, 20].

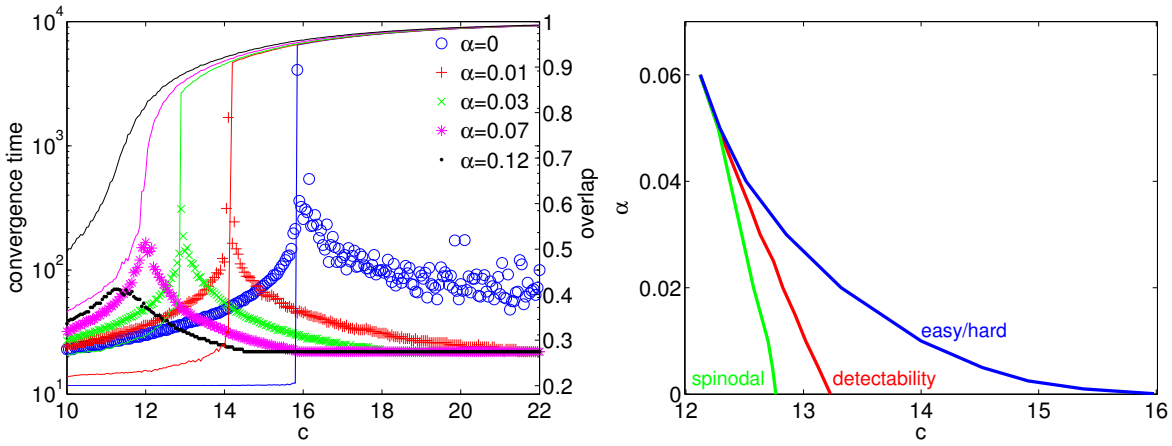

FIG. 3: Left, overlap and BP convergence time in the planted 5-coloring problem as a function of the average degree c for various values of α. Right, the three transitions described in the text: the easy/hard or Kesten-Stigum transition where the factorized fixed point becomes stable (blue), the lower spinodal transition where the accurate fixed point disappears (green), and the detectability transition where the Bethe free energies of these fixed points cross (red). The hard but detectable regime, where BP with random initial messages does no better than chance but exhaustive search would succeed, is between the red and blue lines. All three transitions meet at a critical point, beyond which the overlap is a smooth function of c and α. Here $n = 10^5$.

On the right-hand side of Fig. 1, we show experimental results with $k = 10$. Here the easy/hard transition persists for sufficiently small α, with a discontinuity in the overlap and a diverging convergence time. At a critical value of α, the transition disappears, and the overlap becomes a smooth function of ϵ; beyond that point the convergence time has a smooth peak but does not diverge. Thus there is a line of discontinuities, ending in a second-order phase transition at a critical point. We show this line in the (α, ϵ)-plane in Fig. 2. On the left, we see the discontinuity in the overlap, and on the right we see that the convergence time diverges along this line.

Note that the authors of [20] also predicted the survival of the easy/hard discontinuity in the assortative case. Their approximate computation, however, overestimates the strength of the phase transition, and misplaces its position. In particular, it predicts the discontinuity for all $k > 2$, whereas it holds only for $k > 4$.

The full physical picture of what happens to the "hard but detectable" regime in the semisupervised case, and to the spinodal and detectability transitions that define it, is very interesting. To explain it in detail we focus on the disassortative case, and specifically the case of planted graph coloring where $c_{\text{in}} = 0$. The situation for the assortative case is qualitatively similar, but for graph coloring the discontinuity in the overlap is very strong and appears for any $k > 3$, making these phenomena easier to see numerically.

Fig. 3, on the left, shows the overlap and convergence time of BP for $k = 5$ colors. In the unsupervised case $\alpha = 0$, there are a total of three transitions as we decrease c (making the problem of recovering the planted coloring harder). The overlap jumps at the easy/hard spinodal transition, where the factorized fixed point becomes stable: this occurs at $c = (k-1)^2$. At the lower spinodal transition, the accurate fixed point disappears. In between these two spinodal transitions, both fixed points exist. Their Bethe free energies cross at the detectability transition: below this point, even a Bayesian algorithm with the luxury of exhaustive search would do no better than chance. Thus the "hard but detectable" regime lies in between the detectability and easy/hard transitions [22].

On the right of Fig. 3, we plot the two spinodal transitions, and the detectability transition in between them, in the (c, α)-plane. We see that these transitions persist up to a critical value of $\alpha \approx 0.06$. At that point, all three meet at a second-order phase transition, beyond which the overlap is a smooth function. The very same picture arises in the two related problems mentioned in the introduction, namely the glass transition with random pinning and BP-guided decimation in random k-SAT; see e.g. Fig. 1 in [24, 25] and Fig. 3 in [27].

Finally, in Fig. 4 we plot the overlap and convergence time for the planted 5-coloring problem in the (c, α)-plane. Analogous to the assortative case in Fig. 2, but more visibly, there is a line of discontinuities in the overlap along which the convergence time diverges; the height of the discontinuity decreases until we reach the critical point.

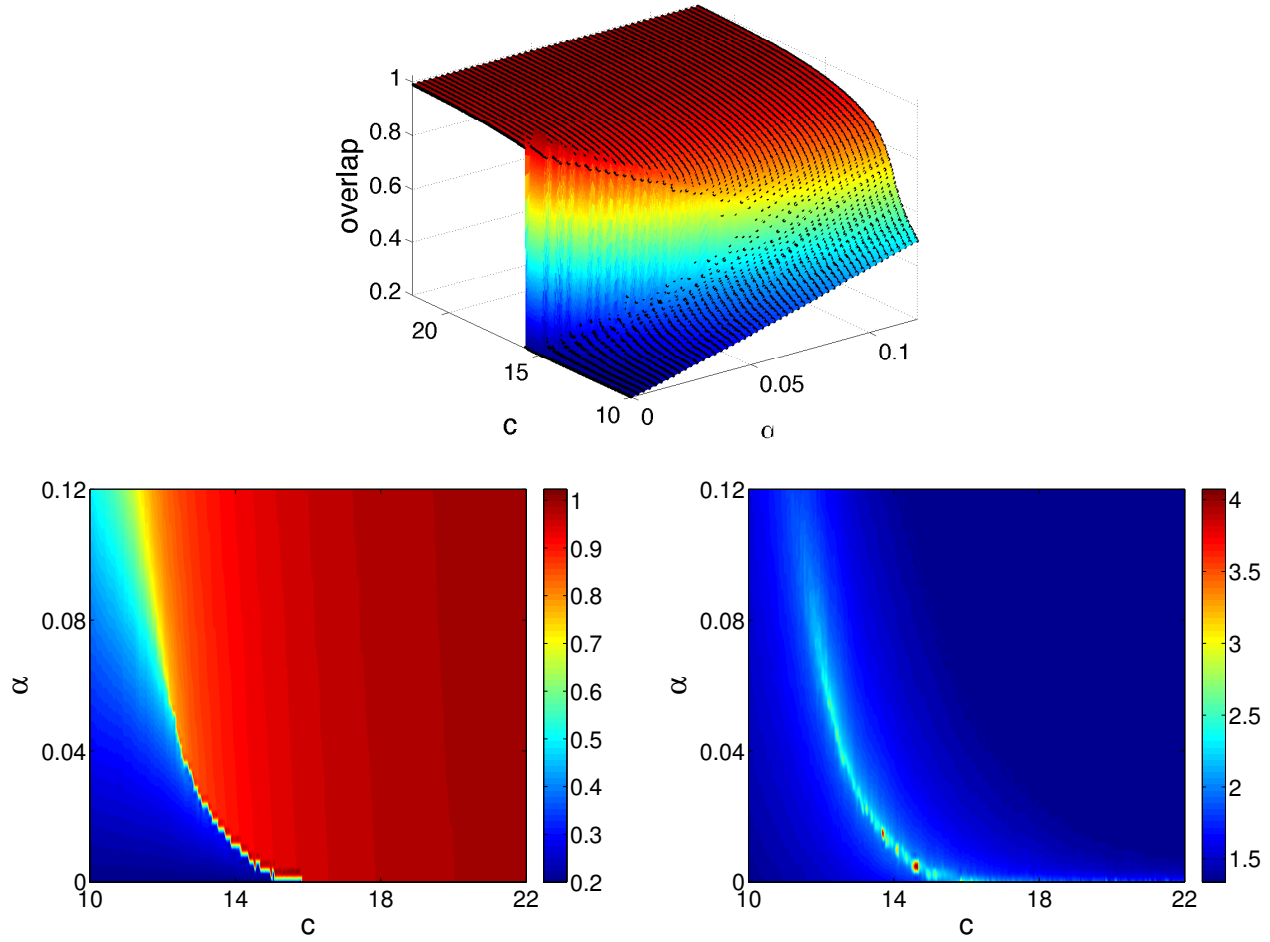

FIG. 4: Overlap (top and left) and convergence time (right) as a function of the average degree c and the fraction of known labels α for the planted 5-coloring problem on networks with $n = 10^5$. The height of the discontinuity decreases until we reach the critical point. The convergence time diverges along the discontinuity. Compare Fig. 2 for the assortative case.

IV. RESULTS ON REAL-WORLD NETWORKS

In this section we study semisupervised learning in real-world networks. Real-world networks are of course not generated by the stochastic block model; however, the block model can often achieve high accuracy for community detection, if its parameters are fit to the network.

To explore the effect of semisupervised learning, we set the parameters q_a and c_{ab} in two different ways. In the first way, the algorithm is given the best possible values of these parameters in advance, as determined by the ground truth labels: this is cheating, but it separates the effect of being given αn node labels from the process of learning the parameters. In the second (more realistic) way, the algorithm uses the expectation-maximization (EM) algorithm of [8, 9], which minimizes the free energy. As discussed in Section II, we initialize the EM algorithm with parameters estimated from edges where both endpoints have known labels, if any.

We test two real networks, namely a network of political blogs [32] and Zachary's karate club network. The blog network is composed of 1222 blogs and links between them that were active during the 2004 US elections; human curators labeled each blog as liberal or conservative. In Fig. 5 we plot the overlap between the inferred labels and the ground truth labels, with multiple independent runs of BP with different initial labels. This network is known not to be well-modeled by the stochastic block model, since the highest-likelihood SBM splits the nodes into a core-periphery structure with high-degree nodes in the core and low-degree nodes outside it, instead of dividing the network along political lines [10]. Indeed, as the top left panel shows, even given the correct parameters, BP often falls into a core-periphery structure instead of the correct one. However, once α is sufficiently large, we move into the basin of attraction of the correct division.

On the top right of Fig. 5, we see a similar transition, but now because the EM algorithm succeeds in learning the correct parameters. There are two fixed points of the learning process in parameter space, corresponding to the high/low degree division and the political one. Both of these are local minima of the free energy [33]. As α increases, the correct one becomes the global minimum, and its basin of attraction gets larger, until the fraction of runs that arrive at the correct parameters (and therefore an accurate partition) becomes large.

We show this learning process in the lower panels of Fig. 5. Since there are just two groups, q_1 determines the group sizes, where we break symmetry by taking q_1 to be the smaller group. As α increases, we move from $q_1 = 0.3$ to $q_1 = 0.47$. On the lower right, we see the parameters c_{ab} change from a core-periphery structure with $c_{22} > c_{12} > c_{11}$ to an assortative one with $c_{11} \approx c_{22} > c_{12}$.

Our second example is Zachary's Karate club [34], which represents friendship patterns between the 34 members of a university karate club which split into two factions. As with the blog network, it has two local optima in parameter place, one corresponding to a high/low degree division (which in the unsupervised case has lower free energy) and the other into the two factions [8]. We again do two types of experiments, one where the best parameters q_a, c_{ab} known in advance, and the other where we learn these parameters with the EM algorithm. Our results are showin in Fig 6 and are similar to Fig. 5. As α increases, the overlap improves, in the first case because the known labels push us into the basin of attraction of the correct division, and in the second case because the EM algorithm finds the correct parameters.

As discussed elsewhere [8, 10], the standard stochastic block model performs poorly on these networks in the unsupervised case. It assumes a Poisson degree distribution within each community, and thus tends to divide nodes into groups according to their degree; in contrast, these networks (and many others) have heavy-tailed degree distributions within communities. A better model for these networks is the degree-corrected stochastic block model [10], which achieves a large overlap on the blog network even when no labels are known. We emphasize that our analysis can easily be carried out for the degree-corrected SBM as well, using the BP equations given in [35]. On the other hand, it is interesting to observe how, in the semisupervised case, even the standard SBM succeeds in recognizing the network's structure at a moderate value of α.

V. CONCLUSION AND DISCUSSION

We have studied semisupervised learning in sparse networks with belief propagation and the stochastic block model, focusing on how the detectability and easy/hard transitions depend on the fraction α of known nodes. In agreement with previous work based on a zero-temperature approximation [19, 20], for $k = 2$ groups the detectability transition disappears for $\alpha > 0$. However, for large k where there is a hard but detectable phase in the unsupervised case, the easy/hard transition persists up to a critical value of α, creating a line of discontinuities in the overlap ending in a second-order phase transition.

We found qualitatively similar behavior in two real networks, where the overlap jumps discontinuously at a critical value of α. When the best possible parameters of the block model are known in advance, this happens when the basin of attraction of the correct structure becomes larger; when we learn them with an EM algorithm as in [8, 9], it occurs because the optimal parameters become global minima of the free energy. In particular, even though the standard block model is not a good fit to networks like the blog network or the karate club, where each community has a heavy-tailed degree distributions, we found that at a certain value of α it switches from a core-periphery structure to the correct assortative structure.

It would be very interesting to apply this formalism to *active* learning, where rather than learning the labels of a random set of αn nodes, the algorithm must choose which nodes to explore. One approach to this problem [36] is to explore the node with the largest mutual information between it and the rest of the network, as estimated by Monte Carlo sampling of the Gibbs distribution, or (more efficiently) using belief propagation. We leave this for future work.

Acknowledgments

C.M. and P.Z. were supported by AFOSR and DARPA under grant FA9550-12-1-0432. We are grateful to Florent Krzakala, Elchanan Mossel, and Allan Sly for helpful conversations.

[1] U. Von Luxburg, Statistics and Computing **17**, 395 (2007).
[2] M. E. J. Newman, Physical Review E **74**, 036104 (2006).

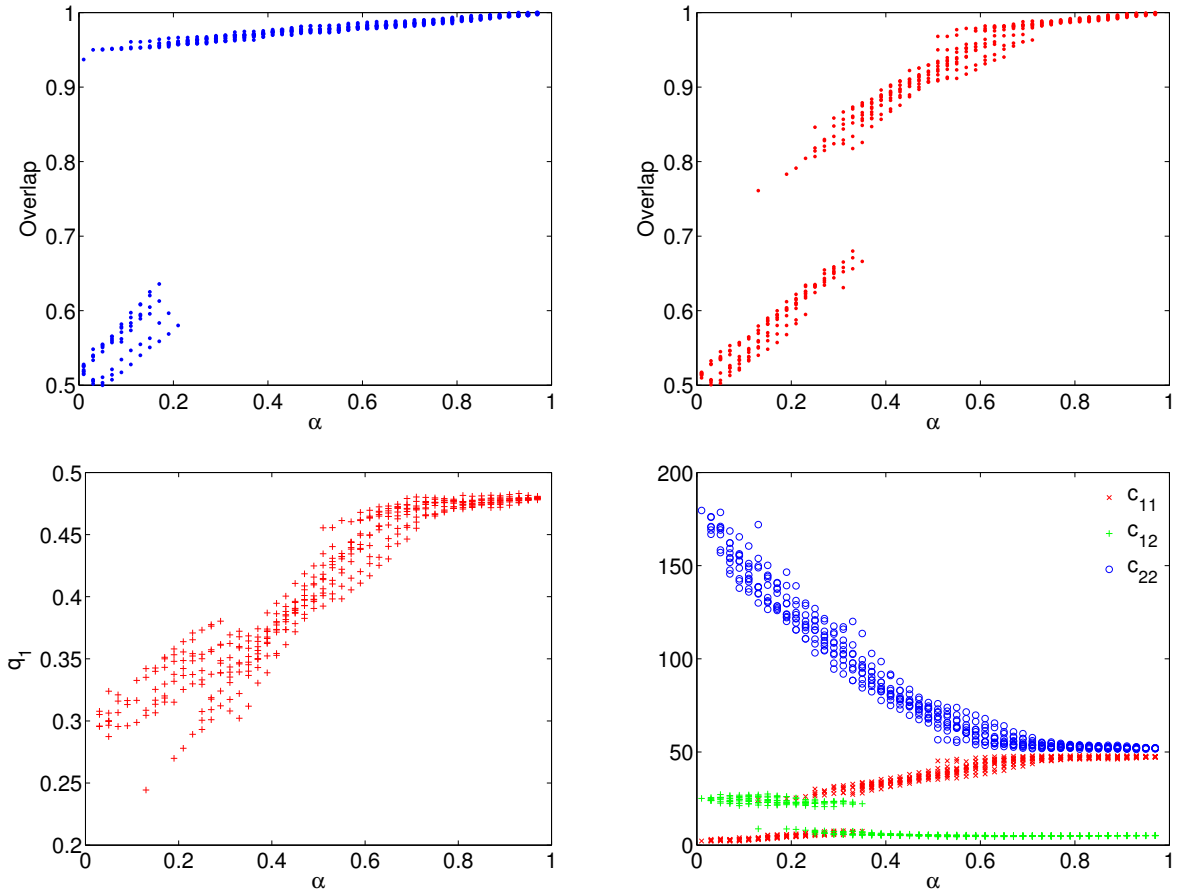

FIG. 5: Semisupervised learning in a network of political blogs [32]. Different points correspond to independent runs with different initial labels. On the top left, the best possible parameters q_a and c_{ab} are given to the algorithm in advance. On the top right, the algorithm learns these parameters using an EM algorithm, seeded by the known labels. The bottom panels show how the learned q_1 and c_{ab} change as α increases, moving from a core-periphery structure where nodes are divided according to high or low degree, to the correct assortative structure where they are divided along political lines.

[3] F. Krzakala, C. Moore, E. Mossel, J. Neeman, A. Sly, L. Zdeborová, and P. Zhang, Proceedings of the National Academy of Sciences **110**, 20935 (2013).
[4] M. E. J. Newman and M. Girvan, Physical Review E **69**, 026113 (2004).
[5] M. E. Newman, Physical Review E **69**, 066133 (2004).
[6] A. Clauset, M. E. Newman, and C. Moore, Physical Review E **70**, 066111 (2004).
[7] J. Duch and A. Arenas, Physical Review E **72**, 027104 (2005).
[8] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Physical Review E **84**, 066106 (2011).
[9] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Physical Review Lett. **107**, 065701 (2011).
[10] B. Karrer and M. E. J. Newman, Physical Review E **83**, 016107 (2011).
[11] V. D. Blondel, J.-L. Guillaume, R. Lambiotte, and E. Lefebvre, Journal of Statistical Mechanics: Theory and Experiment **2008**, P10008 (2008).
[12] M. Rosvall and C. T. Bergstrom, Proceedings of the National Academy of Sciences **105**, 1118 (2008).
[13] S. Fortunato, Physics Reports **486**, 75 (2010).
[14] P. W. Holland, K. B. Laskey, and S. Leinhardt, Social Networks **5**, 109 (1983).
[15] E. Mossel, J. Neeman, and A. Sly, preprint arXiv:1202.1499 (2012).
[16] L. Massoulie, preprint arXiv:1311.3085 (2013).
[17] E. Mossel, J. Neeman, and A. Sly, preprint arXiv:1311.4115 (2013).
[18] O. Chapelle, B. Schölkopf, A. Zien, et al., *Semi-supervised learning*, vol. 2 (MIT press Cambridge, 2006).
[19] A. E. Allahverdyan, G. Ver Steeg, and A. Galstyan, Europhysics Letters **90**, 18002 (2010).
[20] G. V. Steeg, C. Moore, A. Galstyan, and A. E. Allahverdyan, Europhysics Letters (to appear), arxiv.org/abs/1312.0631.
[21] E. Eaton and R. Mansbach, in *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence* (2012).
[22] L. Zdeborová and F. Krzakala, Physical Review E **76**, 031131 (2007).

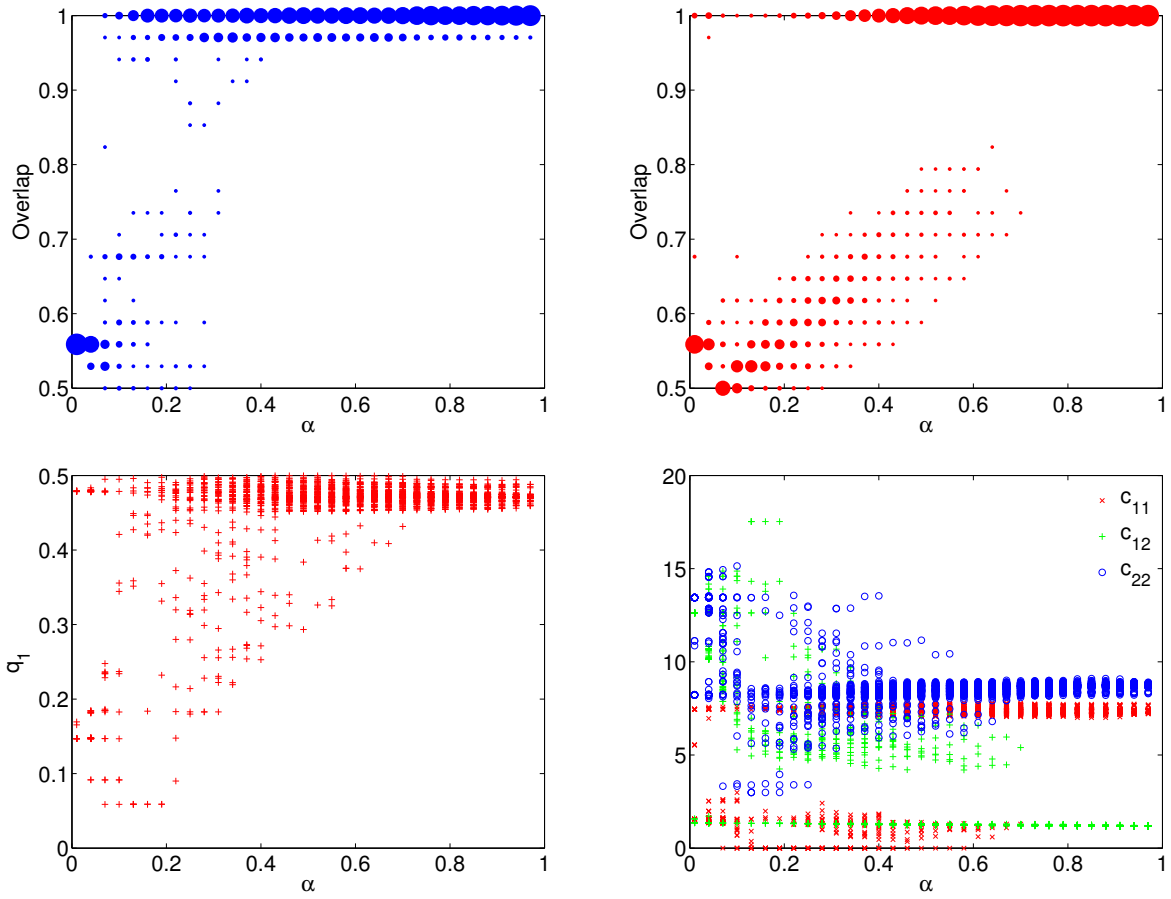

FIG. 6: Semisupervised learning in Zachary's karate club [34], with experiments analogous to Fig. 5: in the upper left, the optimal parameters are given to the algorithm in advance, while in the upper right it learns them with an EM algorithm, giving the parameters shown in the bottom panels. As in the political blog network, the algorithm makes a transition from a core-periphery structure to the correct assortative structure.

[23] M. Mézard and R. Zecchina, Physical Review E **66**, 056126 (2002).
[24] C. Cammarota and G. Biroli, Proceedings of the National Academy of Sciences **109**, 8850 (2012).
[25] C. Cammarota and G. Biroli, The Journal of Chemical Physics **138**, 12A547 (2013).
[26] A. Montanari, F. Ricci-Tersenghi, and G. Semerjian, in *Proceedings of the 45th Allerton Conference* (2007), p. 352359.
[27] F. Ricci-Tersenghi and G. Semerjian, Journal of Statistical Mechanics: Theory and Experiment **2009**, P09001 (2009).
[28] A. Coja-Oghlan, in *Proceedings of the Twenty-Second Annual ACM-SIAM Symposium on Discrete Algorithms* (SIAM, 2011), pp. 957–966.
[29] V. Kanade, E. Mossel, and T. Schramm, preprint arXiv:1404.6325 (2014).
[30] J. Yedidia, W. Freeman, and Y. Weiss, in *International Joint Conference on Artificial Intelligence (IJCAI)* (2001).
[31] M. Mézard and G. Parisi, Eur. Phys. J. B **20**, 217 (2001).
[32] L. A. Adamic and N. Glance, in *Proceedings of the 3rd International Workshop on Link Discovery* (ACM, 2005), pp. 36–43.
[33] P. Zhang, F. Krzakala, J. Reichardt, and L. Zdeborová, Journal of Statistical Mechanics: Theory and Experiment **2012**, P12021 (2012).
[34] W. W. Zachary, Journal of Anthropological Research pp. 452–473 (1977).
[35] X. Yan, J. E. Jensen, F. Krzakala, C. Moore, C. R. Shalizi, L. Zdeborová, P. Zhang, and Y. Zhu, Journal of Statistical Mechanics: Theory and Experiment (to appear), arxiv.org/abs/1207.3994.
[36] C. Moore, X. Yan, Y. Zhu, J.-B. Rouquier, and T. Lane, in *Proceedings of KDD* (2011), pp. 841–849.